

ABSTRACT

The method and system described herein presents a technique for performance analysis that helps users understand the communication behavior of their message passing applications. The method and system described herein may automatically classify individual communication operations and reveal the cause of communication inefficiencies in the application. This classification allows the developer to quickly focus on the culprits of truly inefficient behavior, rather than manually foraging through massive amounts of performance data. Specifically, the method and system described herein trace the message operations of Message Passing Interface (MPI) applications and then classify each individual communication event using a supervised learning technique: decision tree classification. The decision tree may be trained using microbenchmarks that demonstrate both efficient and inefficient communication. Since the method and system described herein adapt to the target system's configuration through these microbenchmarks, they simultaneously automate the performance analysis process and improve classification accuracy. The method and system described herein may improve the accuracy of performance analysis and dramatically reduce the amount of data that users must encounter.

